TRANSACTIONS ON DATA PRIVACY 12 (2019) 1–28

Bootstrap Differential Privacy

Christine M. O'Keefe*, **Anne-Sophie Charest** **

*CSIRO, GPO Box 1700, Canberra ACT 2601 AUSTRALIA.

**Département de mathématiques et de statistique, Université Laval, 1045, avenue de la Médecine, Québec G1V 0A6 CANADA.

E-mail: Christine.O'Keefe@csiro.au, anne-sophie.charest@mat.ulaval.ca

Received 8 January 2018; received in revised form 23 August 2018; accepted 22 October 2018

Abstract. This paper concerns the challenge of protecting confidentiality while making statistically useful data and analytical outputs available for research and policy analysis. In this context, the confidentiality protection measure known as differential privacy is an attractive methodology because of its clear definition and the strong guarantees that it promises. However, concerns about differential privacy include the possibility that in some situations the guarantees may be so strong that statistical usefulness is unacceptably low. In this paper, we propose an example of a relaxation of differential privacy that allows confidentiality protection to be balanced against statistical usefulness. We give a practical illustration of the relaxation implemented as Laplace noise addition for confidentiality protection of contingency and magnitude tables. Tables are amongst the most common types of output produced by national statistical agencies, and these outputs are often protected by noise addition.

Keywords. Differential Privacy; Privacy and Confidentiality; Statistical disclosure control; Statistical disclosure limitation; Privacy-preserving data analysis.

1 Introduction

National statistical agencies and similar data custodians face two competing objectives. On the one hand, an important part of their purpose is to make detailed and accurate data and statistical outputs available to researchers, policy-makers and the community. On the other hand, they must meet their legal and other obligations to protect the confidentiality of the data they collect. These two objectives compete in the sense that releasing data as collected may lead to personal information of individuals being revealed, however releasing only confidentiality-protected data or statistical outputs may reduce their usefulness. There is a large body of research on methods designed to achieve both goals simultaneously with an appropriate balance between them, under the name of *statistical disclosure control*, or *statistical disclosure limitation*, see [6, 7, 18, 23].

In order to determine an appropriate balance between confidentiality protection and statistical usefulness, it is important to be able to measure the extent to which each is achieved. Traditional measures for confidentiality protection have been focussed on estimating re-identification risk, and required assumptions about a potential intruder's additional knowledge and computational resources. Amongst several attempts to provide a more robust definition of confidentiality protection, differential privacy has emerged over recent years

as a strong contender. Informally, the property of differential privacy for an analysis essentially guarantees that adding the data of a single individual to a dataset is unlikely to change the output of that analysis by very much, and hence it should be unlikely that much can be learned about the added individual from the analysis output. Differential privacy is an attractive standard because of its clear definition of privacy and the strong guarantees that it promises.

While differential privacy has had a marked impact on theory and literature in computer science [11], it has had far less impact in the statistical literature and statistical practice. The main concern seems to be that the definition of differential privacy does not mention statistical usefulness at all, so the guarantees may be so strong that analysis outputs are altered to the point where they no longer provide sensible inferences, see for example [15]. In order to address this perceived issue, various relaxations of differential privacy have been proposed with the objective of increasing the statistical usefulness of the outputs. While none of these relaxations has yet been widely adopted in practice, we agree with the opinion that: "it is fruitful to consider various relaxations of differential privacy to gain a deeper understanding of the trade-offs between the strength of the privacy guarantee and the accuracy of the data release mechanism" [16].

This paper proposes and studies a new relaxation of the definition of differential privacy, whereby instead of requiring that analysis outputs do not differ by very much under the addition of *any* individual to the dataset, the condition is required to hold for individuals from a more restricted population. In particular, the relaxation requires that the added individual is again an individual already represented in the dataset. The relaxation has the following features:

- Enables weaker confidentiality protection guarantees than differential privacy.
- Provides a clear description of the weaker confidentiality protection guarantee, thus enabling a consideration of whether the relaxation is acceptable in any given situation.
- Provides a data-dependent definition of the weaker confidentiality protection guarantee, that depends only on the given dataset.
- Allows for the generation of data-dependent analysis outputs with the weaker confidentiality protection guarantee.
- Allows a clear description of the accuracy in one important general class of algorithms. In that case, the accuracy is greater than for differential privacy.
- May be useful in scenarios where the population is unknown, since the relaxation depends only on the dataset, which is known to the data custodian implementing the protection.
- Causes relevant quantities to always be bounded, since they are defined on the known dataset. However, although the space of possibilities is discrete, it may still be impractically large for numerical optimisation.

The new relaxation we propose certainly provides weaker guarantees than those associated with differential privacy, however the guarantees that it provides can be clearly described. We can easily construct a class of examples based on an adaptation of the Laplace mechanism, that may be useful for many complex and flexible models since computing

the sensitivity quantity required for the definition of the Laplace Mechanism is often intractable, either because there simply is no upper bound, or because it is hard to calculate the upper bound analytically. Using this adaptation of the Laplace mechanism also generally increases the statistical usefulness of the outputs, by an amount that can be readily quantified. While we do not suggest that the relaxation replace differential privacy, there may be scenarios in which the weaker confidentiality protection under the relaxation may be sufficiently acceptable, especially given the increase in statistical usefulness.

The rest of this paper is structured as follows. Section 2 provides background definitions and notations, and Section 3 describes our proposed relaxation of differential privacy. The rest of the paper is devoted to discussion and demonstration of implementing the relaxation in practice. In order to keep our discussion practical and useful, we focus on confidentiality protection through the addition of random "noise", since that has traditionally been one of the most widely-used methods in statistical agencies. Further, we focus on the scenario of the release of confidentiality-protected tables, since these include the most in-demand and most common type of output currently released by national statistical agencies. In Section 4 we adapt the method of adding Laplace-distributed noise, show that the adaptation is indeed a relaxation of differential privacy, and in Section 5 we illustrate its use in protecting confidentiality of contingency and magnitude tables. We provide a classification for relaxations of differential privacy, and discuss our relaxation in comparison to others proposed in the literature, in Section 6. The paper concludes with a discussion of ideas for further research.

In summary, the particular contributions of this paper include:

- An example of a relaxation of differential privacy that allows confidentiality protection to be balanced against statistical usefulness.
- A practical illustration of the relaxation implemented as Laplace noise addition for confidentiality protection for contingency and magnitude tables.
- A classification framework for variants of differential privacy.

2 Background and preliminaries

Some necessary background, definitions and preliminaries are reviewed in this section. For a useful introduction to the problems and techniques of differential privacy, see [11] and its references.

2.1 Terminology

The differential privacy framework assumes the scenario of a trusted data custodian agency holding a dataset on a secure server, and releasing outputs of statistical analyses on the dataset requested by users [10].

We model a *dataset* as a multiset of rows, where each row contains data corresponding to a single individual. Let $\mathcal{U}$ denote the collection of all possible dataset rows. We denote by $\mathcal{U}^n$ the collection of all datasets of size n with rows in $\mathcal{U}$. We remark that the available information about $\mathcal{U}$ can range from full information to very little information in different scenarios.

The term *output of statistical analyses* is intended to be interpreted quite widely, and can include statistics such as mean or median, counts or contingency tables, parameters determined during model fitting, outputs of hypothesis tests, and even entire (synthetic)

datasets. In all cases we model a statistical analysis on a dataset x as a function f on datasets with output in some range Range(f). Often the output is a real-valued vector, so that $f(x) = (f_1(x), \ldots, f_k(x))$ and Range$(f) \subseteq \mathbb{R}^k$ for some k, where $\mathbb{R}$ denotes the real numbers.

Informally, a *randomized mechanism* $\mathcal{M}$ is an algorithm that takes as its inputs a dataset x and some random numbers, and produces an output in some range. (For a formal definition, see [11, Definition 2.2].) For example, given a dataset of incomes of individuals in an employment survey in some country, an algorithm that calculates the average income and adds an amount of noise randomly chosen from some distribution with parameter depending on the distribution of incomes in the whole population of that country, would be a randomized mechanism.

2.2 Differential privacy

The concept of a randomized mechanism enables a framework for protecting confidentiality for individuals in a dataset. The intuition is that replacing the data of a single individual with the data of another single individual in a dataset should be unlikely to change the output of a statistical analysis on that dataset by very much, and hence it should be unlikely that much can be learned from the output about either individual involved in the replacement. This intuition is embodied in the definition of differential privacy, guaranteeing that a randomized mechanism gives similar outputs when applied to similar input datasets.

Two datasets $x, y \in \mathcal{U}^n$ are *neighbours* if y can be obtained from x by replacing a row of x with a row from $\mathcal{U}$. Thus neighbouring datasets always have the same size as multisets.

Definition 1. [Differential privacy [8, 10]] A randomized mechanism $\mathcal{M}$ satisfies *ε-differential privacy* if for all neighbouring datasets $x, y \in \mathcal{U}^n$, and all measurable $S \subseteq$ Range(M), we have:

$$\Pr(\mathcal{M}(x) \in S) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(y) \in S). \tag{1}$$

Since Equation (1) is required to hold for all neighbouring datasets $x, y \in \mathcal{U}^n$, it can be written equivalently as:

$$e^{-\varepsilon} \leq \frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} \leq e^{\varepsilon}.$$

We call Equation (1) the *differential privacy condition*. It holds trivially when $x = y$. Note that when ε is small, then e^{ε} is approximately equal to $(1+\varepsilon)$, so that the differential privacy condition is essentially that the difference between $\Pr(\mathcal{M}(x) \in S)$ and $\Pr(\mathcal{M}(y) \in S)$ is at most approximately $\varepsilon \times \Pr(\mathcal{M}(y) \in S)$.

In order to protect the output $f(x)$ of applying a function f to a dataset x (where f could be the identity function but is typically some statistical analysis), one traditional approach has been to add some random noise to $f(x)$. Under the differential privacy approach, one constructs a randomized mechanism that satisfies the property in Equation (1) of the definition of differential privacy, and releases $\mathcal{M}(x)$ instead of $f(x)$. While Equation (1) ensures the privacy guarantee, it is also necessary for the data custodian to separately verify that the value $\mathcal{M}(x)$ released by the mechanism is an acceptable approximation to the true value of $f(x)$, and that the value of ε is sufficiently small that the privacy guarantee is meaningful.

2.3 Achieving differential privacy

There have been many differentially private randomized mechanisms proposed in the literature, see [11] and its references for a useful introduction to the most important examples

to date. In this section we describe a broadly applicable randomized mechanism satisfying differential privacy, namely, the Laplace mechanism.

Given a dataset, and a vector-valued function on datasets, the Laplace mechanism computes the output of the function on the dataset and adds a vector of noise values randomly drawn from a Laplace distribution. The actual distribution used is calibrated by a quantity known as the sensitivity of the function over the collection of all possible rows, designed to capture the magnitude by which a single individual's data can change the output of the function in the worst case.

Definition 2. Let f be a function on datasets with output in $\mathbb{R}^k$. The ℓ_1-*sensitivity* of f is the quantity:

$$\Delta f = \max \|f(x) - f(y)\|_1 \tag{2}$$

where the maximum is taken over all neighbouring datasets $x, y \in \mathcal{U}^n$, and $\| \ \|_1$ is the ℓ_1 norm on $\mathbb{R}^k$ defined by

$$\|x\|_1 = \sum_{i=1}^{k} |x_i|.$$

Sometimes the ℓ_1 is dropped when it is understood.

The *Laplace Distribution* $\mathrm{Lap}(b)$ with scale b is the distribution with probability density function:

$$\mathrm{Lap}(z \mid b) = \frac{1}{2b} \exp\left(\frac{|z|}{b}\right).$$

This distribution has mean zero and variance $\sigma^2 = 2b^2$.

Definition 3. [Laplace Mechanism [10]] Given any function f on datasets $x \in \mathcal{U}^n$ with output in $\mathbb{R}^k$, and a parameter ε, the *Laplace mechanism* (with parameter ε) is:

$$\begin{aligned} \mathcal{L}(x) &= f(x) + (Y_1, \ldots, Y_k) \\ &= (f_1(x) + Y_1, \ldots, f_k(x) + Y_k) \end{aligned}$$

where $Y_1, \ldots, Y_k$ are independent and identically distributed (i.i.d.) random variables drawn from $\mathrm{Lap}(\Delta f/\varepsilon)$.

It is straightforward to show that the Laplace mechanism (with parameter ε) satisfies ε-differential privacy, see [11, Theorem 3.6]. The parameter ε is normally understood when referring to the Laplace mechanism, and so is usually omitted.

It is also straightforward to use properties of the Laplace distribution to prove the following probabilistic bound on the accuracy of the Laplace mechanism.

Theorem 4. *[11, Theorem 3.8] Let f be a function on datasets $x \in \mathcal{U}^n$ with output in $\mathbb{R}^k$, and let $\varepsilon > 0$ be a parameter. Given the Laplace mechanism $\mathcal{L}(x)$ and for $\gamma \in (0, 1]$, we have:*

$$Pr\left(\|\mathcal{L}(x) - f(x)\|_\infty \leq \ln\left(\frac{k}{\gamma}\right)\left(\frac{\Delta f}{\varepsilon}\right)\right) \geq \gamma$$

where

$$\|x\|_\infty = \max_{i=1,\ldots,k} x_i$$

*and Δf is the ℓ_1-*sensitivity *of f.*

Theorem 4 implies that for each $i = 1, \ldots, k$, the probability that the ith perturbed value $\mathcal{L}_i(x)$ lies in an interval of width $2\ln(k/\gamma)(\Delta f/\varepsilon)$ centred at the value $f_i(x)$ is at least γ. If we interpret the interval endpoints as bounds on the difference between the observed value $f_i(x)$ and the perturbed value $\mathcal{L}_i(x)$, then they provide a bound on the error due to the perturbation. For that reason, we call the quantity

$$\left(\ln\left(\frac{k}{\gamma}\right)\left(\frac{\Delta f}{\varepsilon}\right)\right)^{-1} \tag{3}$$

the γ-*accuracy* of $\mathcal{L}$, so that greater accuracy corresponds to an interval of smaller width, and a consequently smaller upper bound on the error due to the perturbation. Since error bounds are of interest to analysts seeking to make inferences based on the perturbed data, we interpret the γ-accuracy as one type of utility measure. Note that a higher value for the ℓ_1-sensitivity of a function f will result in lower accuracy under the Laplace mechanism.

3 Bootstrap differential privacy

There have been several examples in the literature of relaxing the definition of differential privacy in order to weaken the confidentiality protection guarantee in the hope of improving statistical usefulness of outputs, see for example [1, 3, 9]. Some of these approaches are reviewed in Section 6, and we are proposing another. In our relaxation of the definition of differential privacy, we consider restricting the set of neighbouring datasets over which the differential privacy condition Equation (1) is required to hold.

This type of restriction is motivated essentially by two observations. First, suppose that the universe of possible data sets is not known, contrary to the assumption generally made in the case of differential privacy. If there is very little information available about the collection $\mathcal{U}$ of possible rows, then it may be very difficult to determine the impact of replacing a row in the data with any other row in $\mathcal{U}$, and thus to develop provably differentially private randomized mechanisms. A classic example would be that of salary. While we know that salary is a real number, it is not easily capped. Second, in the Laplace mechanism the scale of the noise must be sufficient to ensure the differential privacy condition in the worst case, that is, even for possible dataset rows that are extremely rare or have values that are outliers with respect to any possible values. Since sensitivity is computed on a worst-case scenario, in many situations its value may be quite high, and may lead to unacceptably low mechanism accuracy.

In contrast, everything about an observed dataset D is known to the data custodian implementing the protection, including the presence or absence of rare or outlying individuals. So in theory the two issues outlined above do not arise. In addition, in many situations such as statistical sampling theory and the bootstrap method, the set of individuals in an observed dataset is assumed to be reasonably representative of the collection $\mathcal{U}$ of all possible rows. In that case, the observed dataset D could be used to estimate the impact of replacing a row in the dataset with any other row in $\mathcal{U}$ in order to develop differentially private mechanisms. Given these observations, it may not be unreasonable to propose a relaxation of differential privacy that depends only on the observed dataset. Such a relaxation may be useful in some range of scenarios.

3.1 Definition

Under the relaxation we propose, instead of considering the impact of replacing a row of a dataset D with a row of $\mathcal{U}$ when developing provably private randomized mechanisms, we consider the impact of replacing a row of D with a row of the restricted and known set D itself. In that way, neighbouring datasets D_1, D_2 will differ in a row that must be a row of D. Thus, in the definition of neighbours, instead of allowing the rows of D_1 and D_2 to vary over the full collection of rows in $\mathcal{U}$, we allow them to vary over only the (generally smaller) collection of rows in D. Given a dataset D, the collection of such datasets D_1 and D_2 underpins the bootstrap method [12, 13, 14]. It is this connection, and the notion that the properties of a sample can represent properties of the population, that gives our proposed relaxation its name.

Definition 5. Let D be a dataset of size n. A randomized mechanism $\mathcal{M}$ satisfies *ε-bootstrap differential privacy for D* if for all neighbouring datasets $D_1, D_2 \in D^n$, and all measurable $S \subseteq \text{Range}(\mathcal{M})$, we have:

$$\Pr(\mathcal{M}(D_1) \in S) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D_2) \in S). \tag{4}$$

We remark that the main difference between the definition of bootstrap differential privacy (Definition (5)) in comparison with that of differential privacy (Definition (1)) is that the condition is required to hold over all neighbouring datasets in the subset D^n of $\mathcal{U}^n$, rather than all neighbouring datasets in $\mathcal{U}^n$. This is emphasised by the qualifier "for D" in the definition, though this can be omitted if it is understood.

However, there is an additional difference between bootstrap differential privacy and traditional differential privacy, since the privacy protection is necessarily weaker. The definition of bootstrap differential privacy allows the leakage of information about the dataset D. For example, the mechanism $\mathcal{M}$ that, for any dataset D, always outputs the whole set of distinct rows in D, satisfies the definition of bootstrap differential privacy and clearly leaks the set of distinct rows in D. An example of the potential for leakage of less information about D is discussed in Section 4.1. The data custodian must therefore also separately verify that the leakage of any additional information is acceptable in the given scenario. This would involve understanding the impact of the leakage of such additional information about D on the privacy protection offered by bootstrap differential privacy, and the consequent impact on the property of plausible deniability. We remark that there is a similar (though perhaps not as concerning) situation in the case of (ε, δ)-differential privacy, see the paragraph immediately following [11, Definition 2.4], where mechanism with output a randomly selected small number of rows of the input dataset satisfies (ε, δ)-differential privacy.

Although there exist such examples of differentially private or bootstrap differentially private mechanisms that leak an unacceptable amount of information about the dataset, they would be extremely unlikely to be used in practice - precisely because they leak too much information. Many examples of mechanisms that leak a lot of information in fact do not have a random component, however it is known that randomization is essential to privacy protection, see [11, Section 2.3]. A full theoretical investigation of information leakage in the general case is beyond the scope of this paper.

Finally, just as in the case of differential privacy, the data custodian must separately verify that the mechanism releases an acceptable approximation $\mathcal{M}(D)$ to the true value of $f(D)$, and that the value of ε is sufficiently small that the privacy guarantee is meaningful.

3.2 Privacy guarantees

3.2.1 Plausible deniability

We extend to our case of bootstrap differential privacy, the traditional explanation of the differential privacy guarantee as plausible deniability and interpretation of the parameter ε as privacy loss, see [11, Section 2.3]. In differential privacy, the property of plausible deniability is enjoyed by all individuals, whether they belong to the dataset in question or not. However, in bootstrap differential privacy the property is only guaranteed for individuals represented in the dataset, or with data identical to another individual represented in the dataset.

To be precise, let f be a function on datasets, let D be a fixed dataset, and let $\mathcal{M}$ be a bootstrap differentially private mechanism for D. Suppose that the output $s = \mathcal{M}(D)$ is released, for some value of ε. We consider two cases, according as whether an individual Alice's data is or is not included in the dataset.

First, let d_1 be a row of D, containing Alice's data. Suppose that an intruder guesses that Alice's data is d_1. Alice wishes to deny that d_1 is her row of data, by claiming that instead her row is another row d_2 of D. Alice can choose any row d_2 of D for this purpose, even $d_2 = d_1$. Let $D_2 = (D \setminus d_1) \cup d_2$, and note that D_2 is a bootstrap neighbour of D. Now Equation (4) in the definition of bootstrap differential privacy gives

$$\Pr(\mathcal{M}(D_2) = s) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D) = s).$$

Alice's claim cannot be refuted within the guarantee provided by the parameter ε, since the released value in the case that her row is d_2 cannot be distinguished from the released value in the case that her row is d_1, up to the factor $\exp(\varepsilon)$ as promised by bootstrap differential privacy.

Next, suppose that Alice's data d is not in $D = (d_1, \ldots, d_n)$, and suppose that an intruder knows or guesses Alice's data d. For each $i = 1, \ldots, n$ let $D_i = (D \setminus d_i) \cup d$. If Alice's data d is distinct from each of $d_1, \ldots, d_n$, then D_i is not a bootstrap neighbour of D for any $i = 1, \ldots, n$. In this case, Equation (4) does not hold and Alice cannot plausibly deny that her data are not in D, that is, she cannot plausibly claim that her data are in D. On the other hand, it is possible that $d = d_j$ for some value j (that is, Alice's data row happens to be the same as an individual represented in the dataset). In that case, D_j is a bootstrap neighbour of each $D_i = (D \setminus d_i) \cup d$, and Equation (4) holds for each pair $(D_i, D_j$. In this case, Alice can plausibly deny that her data are not in D, that is, she can plausibly claim that her data are in D.

3.2.2 Outlying observations

To further explore the privacy guarantees provided by this relaxation of differential privacy to bootstrap differential privacy, we consider the impact of outlying observations on analyses of datasets.

Suppose we have a dataset D with the single variable *income* and a bootstrap differentially private mechanism $\mathcal{M}$ on datasets that releases the (privacy-protected) value of the average income of the individuals in the dataset. Suppose the individuals in the dataset come from a population, and let Bill Gates be an individual with income value very much larger than anyone else in the population. We now discuss the impact of replacing an individual in D with another individual (with a different row) in the dataset D, separating the discussion into the two cases that Bill is or is not in the dataset D.

Suppose first that Bill is *in* the dataset D. Under $\mathcal{M}$ the mechanism response would be almost the same if Bill were replaced by a different individual in D (so that Bill's income would be missing from D), or if another individual were replaced by Bill (so that there would be two copies of Bill's income in D). Therefore, any individual in D would be able to plausibly claim that their income was as large as any other person in the dataset, including Bill. Any individual not in D would only be able to make the same claim plausibly if their income was the same as one of the individuals in D, including Bill. Thus, in this case, only individuals with income equal to one of the income values in D (including Bill's income value) enjoy the property of plausible deniability with regard to their income value.

On the other hand, if Bill is *not in* the dataset D, then we can make no statement about how much the mechanism response might change if one of the individuals were replaced by Bill to the dataset. (Note that this scenario should not be confused with the prior situation in which Bill is in the original dataset, then a second copy of Bill's row is added.) Any individual in D would be able to plausibly claim that their income was as large as any other person in the dataset, which does not include Bill. Any individual not in D would only be able to make the same claim plausibly if their income was the same as one of the individuals in D, which does not include Bill. Again, in this case, only individuals with income equal to one of the income values in D (thus not including Bill's income value) enjoy the property of plausible deniability with regard to their income value.

If the mechanism $\mathcal{M}$ were actually differentially private, then any individual either in or out of the dataset D can plausibly claim to have any possible value of income, regardless of whether Bill is in the dataset or not. The plausible deniability of bootstrap differential privacy is clearly less than that of differential privacy, as individuals can only plausibly claim to have an income value that is already in the dataset. This reduced protection may be acceptable in some situations. One such situation might be that the dataset D adequately represents the population - for example, it is a small sample of a census file, or it is known to include all likely values of the relevant variables. In other situations there may be auxiliary information that reduces the scope of an individual's plausible deniability, for example, lifestyle choices may indicate broad limits of income values.

Of course there is an additional challenge if Bill is not initially in the dataset D, but replaces another individual in it after the bootstrap differentially private mechanism has already released the value of average income. If Bill joins D (replacing another individual), then the true average income will change significantly, and the mechanism response may also change significantly, thus revealing information about Bill's income. In addition, since only some individuals in the population have any plausible deniability, it is possible that challenging an individual that does not enjoy such plausible deniability actually reveals the lack of plausible deniability, thus revealing that their income value is different from any income value in the dataset. Bootstrap differential privacy may not be a good choice for privacy protection in cases where it is important to avoid such information leakage.

In summary, the main ways that the privacy protection offered by bootstrap differential privacy is weaker than that of differential privacy are:

- Individuals in the dataset enjoy a reduced scope of plausible deniability. They cannot claim to have any possible data value, but they can claim to have the same value as anyone else in the dataset.

- Outliers in the population that are not in the dataset and with data value different from a value already in the data will not be protected from information leakage if they replace an individual in a dataset.

3.3 Other properties

Bootstrap differential privacy satisfies two important desirable properties of any relaxation of differential privacy, namely strictness and composability.

Result 6. A randomized mechanism $\mathcal{M}$ satisfying ε-differential privacy for some value of the parameter ε also satisfies ε-bootstrap differential privacy for any dataset D, for the same parameter ε. In this case there is no leakage of information about the dataset D. There exist randomized mechanisms satisfying ε-bootstrap differential privacy for some dataset D that do not satisfy ε-differential privacy, for the same parameter ε.

The first statement follows since if the differential privacy condition holds for all neighbouring datasets $D_1, D_2 \in \mathcal{U}^n$, then since $D \subseteq \mathcal{U}$ it also holds for all neighbouring datasets $D_1, D_2 \in D^n$.

For the second statement, the bootstrap Laplace mechanism to be introduced in Section 4.1 will satisfy ε-bootstrap differential privacy for D but not ε-differential privacy. The argument is provided in Section 4.1.

Bootstrap differentially private algorithms for the same dataset D compose according to the following rule.

Result 7. Let D be a dataset, and let $\mathcal{M}_i$ be a ε_i-bootstrap differentially private algorithm for D, for $i = 1, \ldots, k$. Then $(\mathcal{M}_1(x), \ldots, \mathcal{M}_k(x))$ is $\sum_{i=1}^{k} \varepsilon_i$-bootstrap differentially private for D. The information about D that is leaked by $\mathcal{M}$ is the aggregate of the information about D leaked by $\mathcal{M}_1, \ldots, \mathcal{M}_k$.

The proof is essentially the same as the proof of [11, Corollary 3.15].

Finally, we prove another composability result needed later in the paper.

Result 8. Let D be a dataset, and let $D^1, \ldots, D^k$ be a partition of the rows of D. For each $i = 1, \ldots, k$, let $\mathcal{M}_i$ be a randomised mechanism satisfying ε-bootstrap differential privacy for D^i. We define a new mechanism $\mathcal{M}$ on the set of bootstrap samples of D as follows. For a given bootstrap sample B of D, let $(B^1, \ldots, B^k)$ be the partition of B induced by the partition of D, and define

$$\mathcal{M}(B) = (\mathcal{M}_1(B^1), \ldots, \mathcal{M}_k(B^k))$$

Then $\mathcal{M}$ is ε-bootstrap differentially private for D. The information about D that is leaked by $\mathcal{M}$ is the aggregate of the information about D leaked by $\mathcal{M}_1, \ldots, \mathcal{M}_k$.

Proof. The proof is straightforward from the definition of bootstrap differential privacy. Let $D_1, D_2 \in D^n$ be neighbouring datasets, and let $D_1^1, \ldots, D_1^k$ and $D_2^1, \ldots, D_2^k$ be the partitions of D_1, D_2 induced by the partition of D. Suppose without loss of generality that the row of D_1 that has been replaced in order to form D_2 falls into D_1^k and hence also D_2^k. Note that $D_1^i = D_2^i \; \forall i \neq k$. For a measurable subset $S \subseteq \text{Range}(\mathcal{M})$, we have $S = (S_1, \ldots, S_k)$ where

$S_i \subseteq \text{Range}(S_i)$, for $i = 1, \ldots, k$. Then:

$$\begin{aligned}
\Pr(\mathcal{M}(D_1) \in S) &= \Pr\left((\mathcal{M}_1(D_1^1), \ldots, \mathcal{M}_k(D_1^k)) \in S\right) \\
&= \prod_{i=1}^{k-1} \Pr(\mathcal{M}_i(D_1^i) \in S_i) \times \Pr(\mathcal{M}_k(D_1^k) \in S_k) \\
&= \prod_{i=1}^{k-1} \Pr(\mathcal{M}_i(D_2^i) \in S_i) \times \Pr(\mathcal{M}_k(D_1^k) \in S_k) \\
&\leq \prod_{i=1}^{k-1} \Pr(\mathcal{M}_i(D_2^i) \in S_i) \times e^{\varepsilon} \Pr(\mathcal{M}_k(D_2^k) \in S_k) \\
&= e^{\varepsilon} \prod_{i=1}^{k} \Pr(\mathcal{M}_i(D_2^i) \in S_i) \\
&= e^{\varepsilon} \Pr\left((\mathcal{M}_1(D_2^1), \ldots, \mathcal{M}_k(D_2^k)) \in S\right) \\
&= e^{\varepsilon} \Pr(\mathcal{M}(D_2) \in S),
\end{aligned}$$

where we use the independence of the mechanisms $\mathcal{M}_i, i = 1, \ldots, k$ and apply the definition of bootstrap differential privacy for each. □

4 Bootstrap differential privacy through noise addition

The addition of noise to data has traditionally been one of the most widely-used methods for statistical disclosure control in confidential datasets and the outputs of statistical analyses on such datasets, see for example [7, 17, 23]. For that reason, and inspired by the Laplace mechanism, in this section we focus on achieving bootstrap differential privacy via mechanisms that add random noise.

4.1 Bootstrap Laplace mechanism

Adapting the definitions of ℓ_1-sensitivity and the Laplace mechanism in Definitions 2 and 3 respectively gives an example of a bootstrap differentially private mechanism.

Definition 9. [Bootstrap sensitivity] Let D be a dataset, and let f be a function on datasets with output in $\mathbb{R}^k$. The *ℓ_1-bootstrap sensitivity of f for D* is the quantity:

$$\Delta_D f = \max \|f(D_1) - f(D_2)\|_1$$

where the maximum is taken over all neighbouring datasets $D_1, D_2 \in D^n$.

We remark that the only difference between the definition of bootstrap sensitivity in comparison with that of (traditional) sensitivity is that the maximum is taken over all neighbouring datasets in the subset D^n of $\mathcal{U}^n$. Again, this is emphasised by the qualifier "for D" in the definition, though this can be omitted if it is understood.

The notion of bootstrap sensitivity is somewhat reminiscent of local sensitivity [22]. In particular, *local sensitivity* at D is the maximum of $\|f(D_1) - f(D)\|_1$ over the neighbours D_1 of D. However, bootstrap sensitivity differs from local sensitivity in two ways. First, bootstrap sensitivity allows both datasets to vary while local sensitivity fixes one, and second,

bootstrap sensitivity restricts both datasets to be in D^n while local sensitivity has one in D^n and lets the other vary over $\mathcal{U}^n$.

Given a dataset D, the ℓ_1-bootstrap sensitivity captures the magnitude by which the replacement of any row in D by any other row in D can change the value of the function f, in the worst case. Recall that in contrast, the ℓ_1-sensitivity of Definition 2 has to capture such worst-case magnitude for the replacement of any row of $\mathcal{U}$.

Definition 10. [Bootstrap Laplace mechanism] Let D be a dataset. Given any function f on datasets with rows in D and with output in $\mathbb{R}^k$, and a parameter ε, the *bootstrap Laplace mechanism (with parameter ε) for D* is:

$$\begin{aligned}\mathcal{L}(x) &= f(x) + (Y_1, \ldots, Y_k) \\ &= (f_1(x) + Y_1, \ldots, f_k(x) + Y_k)\end{aligned}$$

where Y_i are i.i.d. random variables drawn from $\mathrm{Lap}(\Delta_D f/\varepsilon)$.

We remark that the main difference between the definition of bootstrap Laplace Mechanism in comparison with that of the (traditional) Laplace Mechanism is that the maximum is taken over all neighbouring datasets in the subset D^n of $\mathcal{U}^n$. This is emphasised by the qualifier "for D" in the definition, though this can be omitted if it is understood. The parameter ε is normally understood when referring to the bootstrap Laplace mechanism, and so is usually omitted.

One can easily prove the following result, with an argument analogous to that in [11, Theorem 3.6]:

Result 11. Let D be a dataset and let f be a function on datasets $x \in D^n$ and with output in $\mathbb{R}^k$. Let $\Delta_D f$ be the ℓ_1-bootstrap sensitivity of f. The bootstrap Laplace Mechanism $\mathcal{L}$ with parameter ε satisfies ε-bootstrap differential privacy.

In fact, the bootstrap Laplace mechanism satisfies ε-bootstrap differential privacy but not ε-differential privacy. To see this, consider a single attribute and the function f on datasets that returns the maximum value of that attribute over the individuals in a dataset. Let D be a dataset and suppose that there is an element in the population with value of the attribute greater than the maximum over the individuals in D. The ℓ_1-bootstrap sensitivity of f is less than the ℓ_1-sensitivity of f. The bootstrap Laplace mechanism will satisfy ε-bootstrap differential privacy but not ε-differential privacy.

There are two potential advantages to considering the bootstrap Laplace mechanism instead of the ordinary Laplace mechanism. First, the bootstrap sensitivity of any function f for a dataset D should be less than the sensitivity of f, so the bootstrap Laplace mechanism will require Laplace noise of smaller scale and will thus have higher accuracy (see Section 4.2). Second, since every row of D is known, the bootstrap sensitivity of f for D can in theory be determined, and is also bounded on the collection of datasets with rows in D. If it is not possible to compute the bootstrap sensitivity of f for D directly, or if the problem is computationally infeasible due to the size of the dataset, then it may be necessary to use some discrete optimisation or other algorithm. In any case the search space for bootstrap sensitivity is bounded and would generally be smaller than the search space for sensitivity, although it is not clear that this would make the computation easier in all cases. This situation is different to the case for differential privacy, where in extreme cases sensitivity may be unbounded or it may be hard to calculate analytically.

In addition to separately verifying that the bootstrap Laplace Mechanism gives an acceptable approximation to $f(D)$, the data custodian needs to also verify that any released or

leaked information is acceptable. For example, the bootstrap Laplace mechanism may leak information about the bootstrap sensitivity. In some cases this will be acceptable and in other cases not. (Note in particular that the bootstrap sensitivity is revealed if the mechanism is published.) Consider, as an example, the case of a dataset comprising a single binary variable, and the function that counts the number of 1's. The sensitivity of this function for a dataset D is usually 1, but is 0 in the case that the count is 0 or n. In many practical situations, this does not reveal too much information about the dataset, for two reasons. If the sensitivity is 1, then leaking this value simply reveals that there are both 0s and 1s in the dataset, which would be the expected situation. On the other hand, if the sensitivity is 0, then either everybody or nobody in the dataset has the value 1. Leaking this value would, therefore, be likely to be unacceptable to a data custodian, see [17].

4.2 Accuracy

We can also compute the gain in accuracy that is achieved by the bootstrap Laplace mechanism in comparison to the Laplace mechanism. Recall from Section 2.3 that we can interpret the γ-accuracy as one type of utility measure. This gain in accuracy quantifies the advantage inherent in the reduced privacy protection.

Result 12. Let D be a dataset and let f be a function on datasets $x \in D^n$ and with output in $\mathbb{R}^k$. Let Δf be the ℓ_1-sensitivity of f, and let $\Delta_D f$ be the ℓ_1-bootstrap sensitivity of f. Let $\mathcal{L}$ denote the Laplace mechanism and let $\mathcal{L}_D$ denote the bootstrap Laplace mechanism for D, both with the same parameter ε. Then for $\gamma \in (0, 1)$ the gain in γ-accuracy of $\mathcal{L}_D$ over $\mathcal{L}$ can be represented by the non-negative difference:

$$\left(\ln\left(\frac{k}{\gamma}\right)\left(\frac{\Delta_D f}{\varepsilon}\right)\right)^{-1} - \left(\ln\left(\frac{k}{\gamma}\right)\left(\frac{\Delta f}{\varepsilon}\right)\right)^{-1}$$

or the ratio:

$$\frac{\Delta f}{\Delta_D f}$$

Proof. Since $\Delta f \geq \Delta_D f > 0$, $\varepsilon > 0$, $0 < \gamma < 1$ and $k \geq 1$, then it is straightforward that $\ln(k/\gamma)(\Delta_D f/\varepsilon) \leq \ln(k/\gamma)(\Delta f/\varepsilon)$ and the result follows from the definition of γ-accuracy in Equation (3). □

4.3 Weaknesses and guidance

In this section we further explore and analyse the weaknesses of bootstrap differential privacy in comparison with differential privacy. We provide specific guidance to data custodians on how to decide whether bootstrap differential privacy is acceptable in a given situation). In both cases, we focus in particular on the case of the Laplace mechanism and bootstrap Laplace mechanism.

The main weakness of bootstrap differential privacy, already discussed in this paper, is that the privacy guarantees are weaker than those provided by differential privacy, in particular:

- Privacy guarantees provided to individuals in the dataset are weaker than those provided by differential privacy.

- Privacy guarantees are not provided to individuals not present in the dataset, including individuals who later join the dataset. This is a particular problem for individuals with values that are outliers with respect to the dataset.

In addition, the amount of information about the protection mechanism that can be published is more restricted under bootstrap differential privacy, as follows. Under differential privacy, details of the protection mechanism used to generate the output can be published along with the output, without degrading the privacy guarantees. (See, for example, [11, p88] or [19].) This is helpful for users wishing to take the noise into account in their analysis and inference. Under bootstrap differential privacy, unfortunately publication of details of the mechanism may reveal information about the dataset. In the case of the bootstrap Laplace mechanism, publishing the mechanism would reveal the value of the bootstrap sensitivity, which in turn reveals information about the original dataset. For example, suppose the original dataset is $D = \{1, 2, 3\}$, and we are interested in publishing the sum $f(D)$ of the values in the dataset. Then the bootstrap sensitivity of f is $\Delta_D f = 2$, and the Laplace mechanism with parameter ε operates by adding Laplace noise drawn from the distribution $\text{Lap}(\Delta_D f/\varepsilon)$. An intruder knowing the parameter of the Laplace noise, and the value of ε, can work out the value of the sensitivity. Suppose now that the intruder also knows all observations in the dataset D apart from one, say 1 and 2. With all this information, the intruder can determine that the unknown dataset element is either 0 or 3. The probability of a correct guess (given no further information) is therefore $1/2$. Continuing this thought experiment, suppose the intruder only knows $n-2$ observations as well as the value of the sensitivity. Then still some information about the original dataset is leaked. For example, if the intruder knows the value 1 and that the sensitivity is 2, then it is easy to conclude that the other two observations can not be smaller than -1 or larger than 3.

Therefore, in contrast to the situation under differential privacy, publishing the mechanism used to achieve bootstrap differential privacy is likely to leak information about the dataset. Depending on the context, one may prefer not to release the mechanism used. Even if the mechanism is not released, it is still necessary to verify that the information leaked is within acceptable limits.

Given these weaknesses, situations in which bootstrap differential privacy may be useful include situations where:

- Higher statistical usefulness is considered to justify weaker formal data-centric privacy guarantees.
- The use of other privacy protections such as secure data centres or user registration is considered to justify weaker formal data-centric privacy guarantees.
- The dataset is fixed with respect to participants.
- The population is relatively homogeneous, with few or no outliers.
- The worst-case scenario in computing the sensitivity is extremely rare.
- Very little is known about the collection of possible dataset rows.
- Calculating the sensitivity (for the Laplace Mechanism) is intractable, either because there simply is no upper bound to it; or it is hard to calculate the upper bound analytically. This can occur in complex or flexible models.
- It is not expected that details of the protection mechanism be released.

Further, in addition to the usual checks that the mechanism releases an acceptable approximation to the true value of $f(x)$, and that the value of ε is sufficiently small that the privacy guarantee is meaningful, under bootstrap differential privacy it is necessary to check that the amount of information leaked is acceptable in the given scenario.

5 Illustrations

In this section we illustrate the use of the bootstrap Laplace mechanism in protecting confidentiality when releasing a count or a total computed on a dataset, as well as releasing a contingency table and a magnitude table. After presenting the methodology in Sections 5.1 and 5.2, we provide a numerical example in Section 5.3. We have chosen to focus on tables since they include the most in-demand and most common type of output currently released by national statistical agencies, and this is likely to remain the case for the foreseeable future, see [4].

5.1 Counts and contingency tables

Suppose that given a dataset D of n rows, we wish to release the number $f(D)$ of rows in D that satisfy a certain property. Given any dataset D_1 with rows in D, replacing a row of D_1 by any row of D, can only change the count $f(D_1)$ by at most 1. In the particular case where the count $f(D) = 0$, so that no row of D satisfies the property, then any dataset D_1 with rows in D will also have $f(D_1) = 0$. Similarly, if $f(D) = n$, so that all rows satisfy the property, all datasets D_1 with rows in D will also have $f(D_1) = n$. Thus, if D has at least one but not all rows with the required property, then the bootstrap sensitivity of f for D is 1, and otherwise it is 0. For example, suppose that $f(D)$ computes the number of odd numbers in the dataset D. Then, its sensitivity will be 1 if $D = \{1, 2, 3\}$, but it will be 0 if $D = \{2, 4, 6\}$ or if $D = \{1, 3, 5\}$.

Thus, the definition of the ε-bootstrap Laplace mechanism for a given observed dataset D depends on D itself as well as the function of interest f. If the dataset of interest is such that $f(D) = 0$ or $f(D) = n$, one can directly release the value of $f(D)$ and still satisfy ε-bootstrap differential privacy for any $\varepsilon > 0$. Otherwise, the mechanism

$$\mathcal{M}(x) = f(x) + Y$$

where Y is a value drawn from Lap$(1/\varepsilon)$, is an ε-bootstrap differentially private mechanism for D.

Now suppose we wish to release a contingency table $f(D)$ constructed on D, where in cell i the value $f_i(D)$ is the number of rows in D satisfying the properties that define the cell. Given any dataset D_1 with rows in D, replacing a row of D_1 by a row of D can only change the count in exactly one cell, since the cells partition the dataset. Since each count can be changed by at most 1, the maximum ℓ_1 difference between two contingency tables from neighbouring datasets will be 1, and thus the bootstrap sensitivity of $f(D)$ is 1. We can therefore generalize the mechanism $\mathcal{M}$ above to a mechanism on contingency tables, such that for each cell $f_i(D)$ of the contingency table,

$$\mathcal{M}_i(x) = f_i(x) + Y_i$$

where Y_i are observations drawn independently from a Lap$(1/\varepsilon)$ distribution. Then $\mathcal{M}$ is a ε-bootstrap differentially private mechanism for D. Note that this mechanism is also

ε-differentially private. Thus, bootstrap differential privacy only offers in this case a small improvement in utility over differential privacy, namely the possibility to publish counts of 0 or n directly. This is because the sensitivity is easily bounded (by the value 1) for contingency tables. Bootstrap differential privacy potentially offers more benefits in the case of large or unbounded sensitivity, as will be discussed in the next section.

5.2 Totals and magnitude tables

A magnitude table differs from a contingency table in that the number of individuals in each cell is replaced by the total of the values of a variable for individuals in that cell.

First we consider the case of designing a bootstrap Laplace mechanism to release a single total $f(D)$ on a dataset D. Two neighbouring datasets D_1 and D_2 with rows in D will only differ in the value of one observation, say the n^{th} one. The total will thus be reduced or increased by the difference between the n^{th} observation in D_1 and the n^{th} observation in D_2. The bootstrap sensitivity of a total is thus the maximum absolute difference between two observations of D, which we denote V. Define a mechanism $\mathcal{M}$ by:

$$\mathcal{M}(x) = f(x) + Y$$

where Y is a value drawn from Lap(V/ε). Then $\mathcal{M}$ is a ε-bootstrap differentially private mechanism for D. Note that the mechanism depends on D through the value of V.

Note also that to satisfy the original differential privacy definition, the bootstrap sensitivity for dataset D, $\Delta_D f = V$, would need to be replaced in the mechanism by Δf, the general sensitivity of the function $f(D)$, which depends on the universe $\mathcal{U}$. The gain in accuracy of using ε-bootstrap differential privacy instead of the original ε-differential privacy to release a total will depend on how much larger Δf is compared to $\Delta_D f$. In particular, if the variable of interest is unbounded, then the sensitivity Δf will be infinite so that there will not exist any differentially private mechanism to output the total, whereas the bootstrap differentially private mechanism will still be valid.

Now, suppose we wish to design a bootstrap Laplace mechanism to release a magnitude table $f(D)$ constructed on D, where in cell i the value $f_i(D)$ is the sum of the values for the rows in D satisfying the properties that define that cell. Let k denote the number of cells in the magnitude table. The bootstrap differential privacy sensitivity will be different for each $f_i(D)$, and equal to V_i the maximum absolute difference between two observations in cell i. To release the entire magnitude table with ε-bootstrap differential privacy, we have at least two options.

(i) Consider the entire magnitude table as a single output. Neighbouring datasets D_1 and D_2 with rows in D will differ in only one cell, so that the sensitivity of the entire magnitude table will be the maximum sensitivity for all of the cells:

$$\Delta_D f = \max \|f(D_1) - f(D_2)\|_1 = \max_{i=1,\ldots,k} V_i = V_{imax}.$$

The mechanism $\mathcal{M}$ which computes for each cell $\mathcal{M}_i(x) = f_i(x) + Y_i$, where Y is a value drawn from Lap(V_{imax}/ε) will be ε-bootstrap differentially private for releasing the magnitude table for D.

(ii) Consider each $f_i(D)$ as a separate output from D. By the composition property of bootstrap differential privacy, we can simply release each total $f_i(D)$ using the bootstrap Laplace mechanism with the sensitivity $\Delta_D f_i = V_i$ and privacy parameter ε/k to obtain an ε-bootstrap differentially-private mechanism for dataset D.

The best approach, in terms of utility of the released magnitude table, will depend on the number k of cells in the table as well as the variability in the sensitivities $\{V_i, i = 1, \ldots, k\}$. Approach (ii) uses a smaller privacy parameter ε/k, which will require adding more noise, but allows for less noise to be added to cells with smaller sensitivities. If one or a few V_i are much larger than the others, it may thus be advantageous to use option (ii).

Note that similar approaches could also be used to satisfy the original differential privacy definition. Under option (i), one would use $\Delta f = V$, which depends on U, instead of $\Delta_D f = V_{imax}$. Just as for the release of a single total, the gain of utility of using bootstrap differential privacy instead of differential privacy will depend on the difference between these two quantities. Under option (ii), the sensitivity used for each cell could be different only if we can *a priori* identify different universes $\mathcal{U}_i$ for the values of the variable of interest in each of the cells. This may however not be easily achieved in a practical situation. Bootstrap differential privacy is thus advantageous in this case.

Note also that a third option is available if one is willing to make a further assumption, namely that cell membership for each observation is public information. This may be a fair assumption in some cases, for example if the dataset is about companies and the cells are formed based on geography and type of business. Under the assumption that cell memberships are public, one can, without any privacy loss, split the observations into k datasets where dataset D_i contains the observations classified into cell i. Then, releasing the total in each of the cells represent independent queries. By Result 8, a mechanism M which uses independently the bootstrap Laplace mechanism for each dataset D_i, with cell-specific bootstrap sensitivity V_i and privacy parameter ε will provide overall ε-bootstrap differential privacy for the whole magnitude table for dataset D. Note that under this assumption, bootstrap differential privacy still provides plausible deniability, but now an individual can only claim to have value equal to any value from among the values in the cell he belongs to, including his, instead of any value among all values in the dataset.

We now illustrate some of these results with a numerical example.

5.3 Numerical example

In this section we illustrate the bootstrap Laplace mechanisms for a magnitude table as described in Section 5.2 on a small real dataset. We also offer some numerical comparisons with the corresponding differentially private Laplace mechanisms. We use $\varepsilon = 1$ for all examples.

We use a dataset D of attributes of rice farms in India, available as a standalone file associated with the `plm` library of the R statistical software package [5, 24]. The dataset contains 1026 observations, with 21 variables. Although the dataset actually comprises a time series of 6 observations on each of only 171 farms, we will treat it as a single observation on each of 1026 different farms for the purpose of this illustration. Even on repeated observations of the same farm, the values can vary quite markedly as event status, size and varieties produced are not constant over time.

For the purposes of our illustration, the farms are classified according to the variable `Status`, with values *Mixed*, *Owner*, and *Share*, and the variable `Varieties`, with values *High*, *Both*, and *Traditional (Trad)*. Table 1 shows, for each cell of this two-way classification, the number n of farms falling into that cell, the total net output `total nout` of rice measured in kilograms, and the maximum output `max` and minimum output `min` of rice of any farm in the cell, again in kilograms. Note that the net output is the value which would be shared in the magnitude table.

Table 1: Two-way classification of rice farms in India by Status and Varieties, showing number n of farms, net output (`nout`) of rice (kg), maximum output (`max`) and minimum output (`min`) of rice (all in kg) for farms in the cell

		Varieties		
Status		**High**	**Both**	**Trad**
Mixed	n	33	7	171
	`total nout`	56,965	11,187	189,528
	`max`	900	3,400	3,200
	`min`	234	800	180
Owner	n	227	41	468
	`total nout`	416,820	76,917	436,757
	`max`	17,610	12,000	8,100
	`min`	82	200	42
Share	n	34	2	43
	`total nout`	58,669	1,105	25,236
	`max`	14,520	705	2,000
	`min`	184	400	100

5.3.1 Publishing a cell total

We first consider the task of publishing one cell total with bootstrap differential privacy. To do so, we simply apply the bootstrap Laplace mechanism. We present here some results for each of the cells in the magnitude table, to illustrate the properties of the mechanism. Note that each cell total will be protected at level ε, so that the overall magnitude table would have protection 9ε if released.

Individual cell sensitivities are given in table 2. Figure 1 shows the distribution of 1000 independent replicates of the bootstrap Laplace mechanism for the total net output of each cell. Totals for cells with smaller bootstrap sensitivity are less perturbed by the mechanism.

Table 2: Bootstrap cell sensitivities for the farm example

	Varieties		
Status	**High**	**Both**	**Trad**
Mixed	8766	2600	3020
Owner	17528	11800	8058
Share	14336	305	1900

Note that by design of the mechanism the average, and median, of the totals within a cell is expected to be the true cell total. This is illustrated in Figure 2 which shows the distributions of the relative bias for the 1000 replicates of the mechanism.

It is interesting to note that cells with the smallest bias are not necessarily those with the

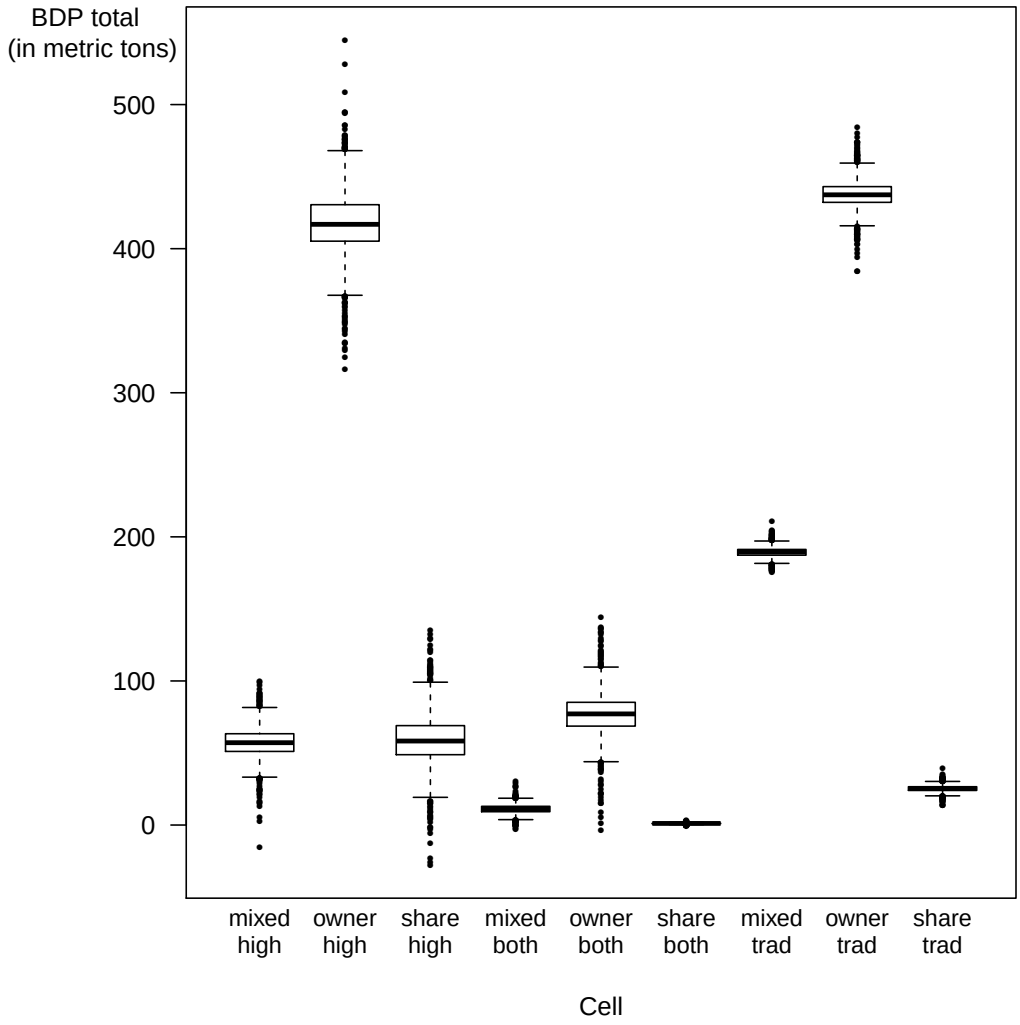

Figure 1: Distributions of 1000 independent replicates of the bootstrap Laplace mechanism for the total net output of each cell. Each cell is perturbed separately by the mechanism with $\varepsilon = 1$ and its own cell specific sensitivity.

smallest sensitivity. In fact, relative bias increases with sensitivity but decreases with size of the true total, and thus indirectly with sample size. Cells where the sensitivity represents a higher proportion of the total will be more perturbed relative to the size of the total, as can be seen from comparing Figure 2 with Table 3 which gives the ratio of the sensitivity to the total, in percentage, for each cell.

It is logical to ask how these results differ from the ones we could obtain with the original differentially private Laplace mechanism. First, note that in general the output of rice from a farm is an unbounded quantity, and thus there would exist no Laplace mechanism to output the cell counts. One may be able to provide ahead of time a maximum value for

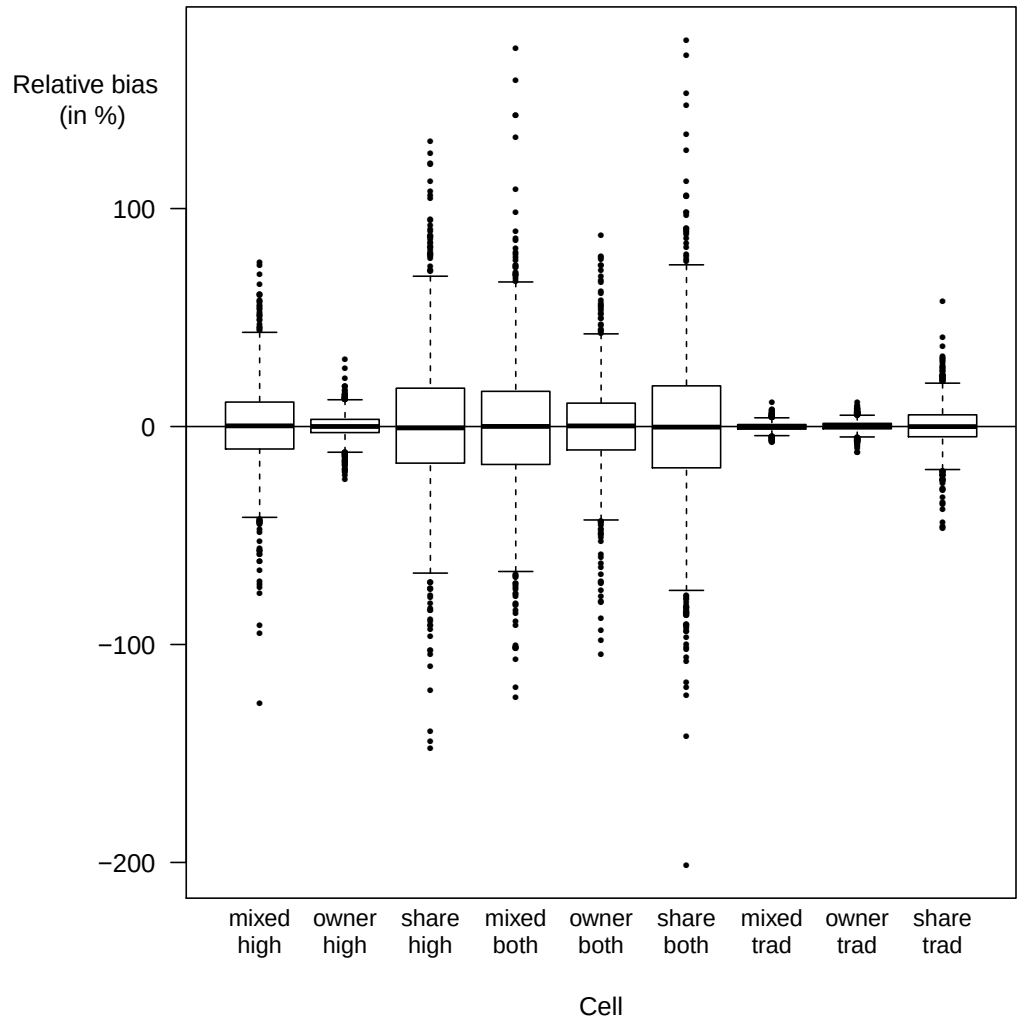

Figure 2: Relative bias (%) of 1000 independent replicates of the bootstrap Laplace mechanism for the total net output of each cell. Each cell is perturbed separately by the mechanism with $\varepsilon = 1$ and its own cell specific sensitivity.

the amount of rice to be produced in a farm, and use this to compute the sensitivity Δf. The smallest plausible value to use in our example would be the maximum output of all observations, namely 17,610. Using this value instead of $\Delta_D f$ would result in a loss of accuracy for many of the nine cells. The ratios of accuracy values

$$\frac{\textit{bootstrap differentially private } \gamma\textit{-accuracy}}{\textit{differentially private } \gamma\textit{-accuracy}} = \frac{\Delta f_i}{\Delta_D f_i}$$

for each of the nine cells under this assumption are given in Table 4. Cells for which the sensitivity is smallest gain most from the use the bootstrap Laplace mechanism.

Table 3: Ratio of bootstrap cell sensitivities to cell totals (percentages %)

	Varieties		
Status	**High**	**Both**	**Trad**
Mixed	15.39	23.24	1.59
Owner	4.21	15.34	1.84
Share	24.44	27.60	7.53

Table 4: Ratios of bootstrap differentially private to original differentially private $\gamma-$accuracy for the farm example if we use the largest observation in the dataset to bound the possible net output on a farm

	Varieties		
Status	**High**	**Both**	**Trad**
Mixed	2.009	6.773	5.831
Owner	1.005	1.492	2.185
Share	1.228	57.738	9.268

Instead of simply using the overall maximum rice output, one could provide a maximum output for the farms in the cell of interest, or estimate that value in a differentially private manner from the observed dataset. In both cases, the bootstrap Laplace mechanism would still provide a gain in utility in comparison with the Laplace mechanism, but in might be much smaller than those in Table 4.

5.3.2 Publishing the entire magnitude table

We have defined earlier two different mechanisms to publish the entire table. In method (i), we use the maximum cell-specific sensitivity to add Laplace noise directly on the vector of totals in the magnitude table. In our example, we thus have $\Delta_D f = 17,528$, whereas $\Delta f = 17,610$ which will lead to a slight improvement in $\gamma-$accuracy. This improvement will be larger in cases where the values vary a lot between cells, but are similar within cells, so that the distance between any two observations will be much larger than the maximum distance between any two observations in the same cell.

In method (ii), we simply use the same mechanism as for publishing a cell, but with $\varepsilon/9$ instead of ε so that the overall magnitude table is generated with privacy parameter ε. This of course results in a reduction in the utility for every published cell compared to when we considered publishing a single cell.

The choice between method (i) and method (ii) will depend on the specifics of the dataset. Figure 3 shows the cell values for 1000 magnitude tables perturbed with each of the two methods for the farms example. We fist note that the noise added here is larger than that required to release a single total, as illustrated in Figure 1. Also, note that the noise added

with method (i) is identical for all of the cells, whereas the noise added with method (ii) is larger for cells with higher bootstrap sensitivity. In this particular example, only one cell (with values share-both) is more perturbed under method (i) than method (ii). Since this cell only contains 2 observations, we may decide to simply not release its total and use method (i) to release the rest of table.

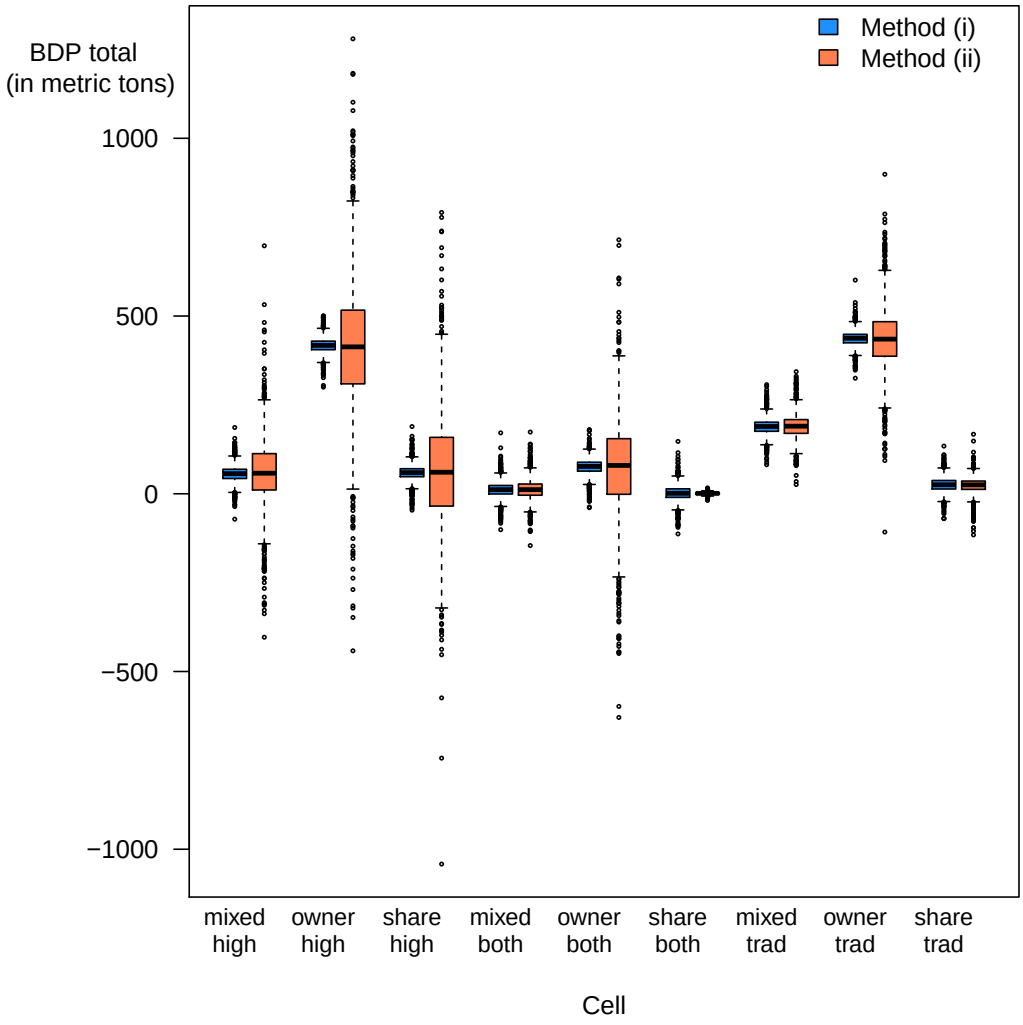

Figure 3: Distribution of 1000 independent replicates of the bootstrap-differentially-private mechanisms (i) and (ii) for the release of the complete magnitude table of the farms example.

6 Variations of differential privacy

Several variations of the definition of differential privacy have been proposed in the literature, now including bootstrap differential privacy. In this section we give a summary of some of these, and give some observations about the relationships between them.

The variations we consider can be classified as one of the following types: relaxation of the bound in the differential privacy condition (*bound relaxation*), restriction of the collection of datasets over which the differential privacy condition is required to hold (*scope restriction*), replacement of the probability distribution in the differential privacy condition (*distribution replacement*), and combinations of these. Bound relaxations and scope restrictions involve relaxations of the differential privacy condition.

6.1 Bound relaxation

The first example of this type of relaxation is (ε, δ)-*differential privacy* [10], with condition, for $D_1, D_2 \in \mathcal{U}^n$:

$$\Pr(\mathcal{M}(D_1) \in S) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D_2) \in S) + \delta.$$

If $\delta = 0$ then (ε, δ)-differential privacy is just ε-differential privacy, and in fact this relaxation is so popular that it is often given as the definition of differential privacy.

A full description of the difference between (ε, δ)-differential privacy and ε-differential privacy is given in [11, Section 2.3].

The second example of a bound relaxation is (ε, γ)-*probabilistic differential privacy* [20], with condition:

$$\Pr\left(\Pr(\mathcal{M}(D_1) \in S) \leq \ e^{\varepsilon}\Pr(\mathcal{M}(D_2) \in S)\right) \geq 1 - \gamma. \tag{5}$$

In probabilistic differential privacy, the overall probability that the differential privacy condition does not hold is bounded.

The outer probability in Equation (5) is calculated with respect to the joint distribution of the possible inputs and outputs of the randomized mechanism $\mathcal{M}$. Since $\mathcal{M}$ determines the distribution of outputs conditional on the inputs, computation requires a distribution over the possible inputs. In the original proposal [20], a uniform distribution over possible inputs is implicit, but one could use any prior distribution, see Section 6.3. The validity of the guarantee in practice will however depend on the validity of the assumed distribution of the dataset.

6.2 Scope restriction

The two examples of this type of relaxation presented here both depend on a given or observed dataset D. Each restricts the collection of pairs of neighbouring datasets from those with rows in $\mathcal{U}$ to those with rows in some smaller collection of rows that depends on D.

In the first example, called ε-*conditional differential privacy* for a given dataset D [2], the differential privacy condition is required to hold:

> for pairs D, D^{-i} where D^{-i} is obtained from D by removing the i^{th} observation.

Conditional differential privacy was defined in order to explore the impact on concepts related to differential privacy on simply removing a row from a dataset, thus its properties were not fully explored. However, one can immediately see that desirable properties such as composability will not hold.

The second example, that of ε-*individual differential privacy* given a dataset D [25], requires that the differential privacy condition hold:

for pairs D, D' where D' is a neighbour of D.

This relaxation is related to the concept of local sensitivity [22]. It is designed to allow the data custodian to calibrate the noise magnitude to an actual dataset, normally resulting in greater analytical accuracy. While the same guarantees as for standard differential privacy hold for individuals under this relaxation, there are no such guarantees for groups of individuals. We remark that calibrating the noise magnitude to an actual dataset may also reveal information about the dataset, similarly to the case of bootstrap differential privacy.

The third example is ε-*bootstrap differential privacy for* D as defined in this paper, for which the differential privacy condition is required to hold:

for all neighbouring datasets $D_1, D_2 \in D^n$.

The scope restriction in ε-conditional differential privacy is stronger than that in ε-bootstrap differential privacy for D, so that there are likely to be more mechanisms satisfying ε-bootstrap differential privacy for D.

6.3 Distribution replacement

The notion of random differential privacy proposes to view the rows of a dataset D as random draws from an unknown distribution $\mathcal{P}$ on $\mathcal{U}$.

Definition 13. [16] A randomized mechanism $\mathcal{M}$ satisfies (ε, γ)-*random differential privacy* if for all $n = 1, 2, \ldots$, for all neighbouring datasets $D_1, D_2 \in \mathcal{U}^n$, and all measurable $S \subseteq \text{Range}(\mathcal{M})$, we have:

$$\Pr\left(\Pr(\mathcal{M}(D_1) \in S) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D_2) \in S)\right) \leq 1 - \gamma$$

where the outer probability is with respect to the $(n+1)$-fold product measure $\mathcal{P}^{n+1}$ on $\mathcal{U}^{n+1}$ (since $|D_1 \cup D_2| = n + 1$).

The relaxation makes use of the same condition as probabilistic differential privacy, namely the condition in Equation (5), however instead of taking the outer probability with respect to the randomized mechanism, it is taken with respect to the $(n+1)$-fold product measure $\mathcal{P}^{n+1}$ on $\mathcal{U}^{n+1}$ (since $|D_1 \cup D_2| = n+1$). This relaxation, then, seeks to relate the differential privacy condition to the generation process that led to the dataset itself.

If we denote the observed values of the random variables $D = (X_1, \ldots, X_n)$ by $d = (x_1, \ldots, x_n)$, then the definition of differential privacy can be informally restated without loss of generality by saying that $\Pr(\mathcal{M}(x_1, \ldots, x_n) \in S)$ is not strongly affected by subtracting a row x_n or adding a row $x_{n+1} \in \mathcal{U}$. Under random differential privacy, this condition is replaced by the informal restatement that $\Pr(\mathcal{M}(x_1, \ldots, x_n) \in S)$ is not strongly affected by subtracting a row x_n or adding a row x_{n+1} randomly drawn from $\mathcal{P}$.

The idea of random differential privacy is somewhat similar to the idea of bootstrap differential privacy, in that there is a restriction on the collection of rows from which a row can be added to a dataset to create a neighbour. In particular, neighbours are defined by removing a row or by adding a row drawn from an (unknown) distribution underlying the population, not any row of the population. In cases where this distribution is very similar to the distribution on the observed dataset, then random differential privacy and bootstrap

differential privacy may give similar results, since the bootstrap distribution is effectively providing an approximation of the distribution underlying the population. As an example, the discussion of releasing counts (in a histogram) provided in [16] shows that zero cells need not be protected under random differential privacy. On the other hand, the confidentiality protection guarantee for random differential privacy may not be so straightforward to describe, given the more technical definition.

6.4 Combinations

Some pairs of relaxations from Sections 6.1, 6.2, and 6.3 can be combined.

For example, the definition of $(\varepsilon, \delta, \gamma)$-random differential privacy [16] combines the random differential privacy relaxation with (ε, δ)-differential privacy relaxation, with corresponding condition, for $D_1, D_2 \in \mathcal{U}^n$

$$\Pr\left(\Pr(\mathcal{M}(D_1) \in S) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D_2) \in S) + \delta(n)\right) \leq 1 - \gamma$$

where δ is a function that decreases faster than any inverse polynomial in n.

We remark that the bootstrap differential privacy relaxation could be combined with other relaxations such as those for (ε, δ)-differential privacy, probabilistic differential privacy, and random differential privacy. These further relaxations can be the subject of future investigations.

7 Conclusion and future research

In this paper we have proposed bootstrap differential privacy as a relaxation of differential privacy. Like differential privacy, bootstrap differential privacy also allows a clear description of the confidentiality protection guarantee, as well as a formal measure of accuracy. The main difference is that bootstrap differential privacy generally provides less confidentiality protection and more accuracy than differential privacy. Another important feature is that any implementation of bootstrap differential privacy will depend only on the observed dataset, and does not rely on any properties of a population. Thus, even if the population is completely unknown, the quantities relevant for bootstrap differential privacy are bounded and can generally be determined unless there is a computational barrier. While we do not suggest that the relaxation replace differential privacy, there may be scenarios in which the weaker confidentiality protection under the relaxation may be sufficiently acceptable, especially given the increase in statistical usefulness.

Given a dataset D and a collection $\mathcal{U}$ of all possible rows, the relationship between bootstrap differential privacy and differential privacy depends somewhat on the degree to which D is representative of $\mathcal{U}$. If D is highly representative of $\mathcal{U}$, such as being a large random sample, then bootstrap differential privacy may be quite similar to differential privacy. In the extreme case that $D = \mathcal{U}$, they will be the same. In the case of the Laplace mechanism, if a function f is very sensitive on rows in $\mathcal{U} \setminus D$ that are rather unlikely to occur, then bootstrap differential privacy requires the addition of less noise and protects against replacement of rows in D with rows that actually occur in the dataset, not rows that are very unlikely to occur.

We have provided a classification for relaxations of differential privacy, and discussed our relaxation in comparison to others proposed in the literature.

We have illustrated the implementation of bootstrap differential privacy in the common scenarios of releasing contingency and magnitude tables with confidentiality protection

provided by the addition of random noise. Our illustration has shown that that is very straightforward to apply bootstrap differential privacy to the release of a contingency table or a magnitude table, and that the reduction in privacy guarantee may be accompanied by a non-negligible increase in statistical usefulness.

Our investigation of bootstrap differential privacy has also raised further questions for future research. Perhaps the most pressing is a full theoretical investigation of the impact of the leakage of any additional information about D, and the consequent impact on the property of plausible deniability.

The idea of bootstrap differential privacy suggests the existence of a family of nested variants of differential privacy that are "in between" bootstrap differential privacy and differential privacy. Such variants would provide a range of confidentiality protection levels understood in terms of the scope of the plausible deniability property. It would be interesting to explore these nested variants in which the differential privacy condition is required to hold for neighbouring datasets with rows in a given defined sub-population, that is smaller than the entire population, but larger than the dataset itself. This might be useful in cases that the population is unknown, or has extremely large outliers in addition to some outliers that are not so extreme. In any given scenario, a sub-population could be chosen amongst the collection of all such sub-populations to provide the most acceptable level of confidentiality protection. As a generalisation, if the acceptable level of confidentiality protection is known in advance, then fake (or synthetic) individuals could be added to the dataset to ensure the desired level of confidentiality protection.

It would be interesting to explore the impact of relaxing to bootstrap differential privacy in the definition of other differentially private mechanisms, such as the exponential mechanism [21]. Other possibilities include exploring other relaxations of bootstrap differential privacy, such as (ε, δ)-bootstrap differential privacy. Of particular interest would be to explore the properties of bootstrap differentially private synthetic datasets, and their suitability in a range of situations.

Another interesting line of investigation would be to explore the potential to exploit the bootstrap method in differential privacy itself. Under an example of the bootstrap method, properties of a quantity of interest on a population $\mathcal{U}$ are estimated by evaluating the quantity on a large number of random samples drawn with replacement from a dataset $D \subseteq \mathcal{U}$. Given D, the bootstrap method may give good estimates for quantities of interest to differential privacy, such as sensitivity as in the definition of the Laplace mechanism. One example of this was discussed in Section 6.3, where the bootstrap distribution provides an approximation of the data generating mechanism required to define and implement random differential privacy.

Acknowledgements

The work of the first author was partially supported by a grant from the Simons Foundation. The work of the second author was partially supported by the Natural Sciences and Engineering Research Council of Canada grant No. RGPIN-435472-2013. Both authors were supported by EPSRC grant no EP/K032208/1.

References

[1] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In *Proceedings of the 26th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 273–282, 2007.

[2] A.-S. Charest and Y. Hou. On the meaning and limits of empirical differential privacy. *Journal of Privacy and Confidentiality*, 7(3):3, 2017.

[3] K. Chaudhuri, C. Monteleoni, and A.D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

[4] J. Chipperfield, D. Gow, and B. Loong. The Australian Bureau of Statistics and releasing frequency tables via a remote server. *Statistical Journal of the IAOS*, 32(1):53–64, 2016.

[5] Y. Croissant, G. Millo, et al. Panel data econometrics in r: The plm package. *Journal of Statistical Software*, 27(2):1–43, 2008.

[6] T. Dalenius. Towards a methodology for statistical disclosure control. *Statistik Tidskrift*, 15:429 – 444, 1977.

[7] G.T. Duncan, M. Elliot, and J.-J. Salazar-Gonzàlez. *Statistical Confidentiality*. Springer, New York, 2011.

[8] C. Dwork. Differential privacy. In M Bugliesi, B Preneel, V Sassone, and I Wegener, editors, *ICALP 2006*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12, Heidelberg, 2006. Springer.

[9] C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*, pages 371–380, 2009.

[10] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *3rd IACR Theory of Cryptography Conference*, pages 265–284, 2006.

[11] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[12] B. Efron. The 1977 Reitz Lecture. Bootstrap methods: another look at the jackknife. *Annals of Statistics*, 7:1–26, 1979.

[13] B. Efron and R. Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1:54–57, 1986.

[14] B. Efron and R.J. Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.

[15] S.E. Fienberg, A. Rinaldo, and X. Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *International Conference on Privacy in Statistical Databases*, pages 187–199. Springer, 2010.

[16] R. Hall, A. Rinaldo, and L. Wasserman. Random differential privacy. *Journal of Privacy and Confidentiality*, 4(2):43–59, 2012.

[17] A. Hundepool, J. Domingo-Ferrer, L. Franconi, S. Giessing, R. Lenz, J. Naylor, E.S. Nordholt, G. Seri, and P-P. DeWolf. Handbook on statistical disclosure control. http://neon.vb.cbs.nl/casc/SDC_Handbook.pdf, 2010. Accessed 23 Jan 2013.

[18] A. Hundepool, J. Domingo-Ferrer, L. Franconi, S. Giessing, E.S. Nordholt, K. Spicer, and P.-P. de Wolf. *Statistical Disclosure Control*. Wiley Series in Survey Methodology. John Wiley & Sons, United Kingdom, 2012.

[19] B.-R. Lin and D. Kifer. Towards a systematic analysis of privacy definitions. *Journal of Privacy and Confidentiality*, 5(2):2, 2014.

[20] A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the IEEE 24th International Conference on Data Engineering ICDE*, pages 277 –286, April 2008.

[21] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium*, pages 94–103. IEEE, 2007.

[22] K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, Pages 75–84, 2007.

[23] C.M. O'Keefe and D.B. Rubin. Individual privacy versus public good: protecting confidentiality in health research. *Statistics in Medicine*, 34:3081–3103, 2015.

[24] R Project for Statistical Computing, n.d. `http://www.r-project.org/`.

[25] J. Soria-Comas, J. Domingo-Ferrer, D. Sanchez and D. Meges. Individual differential privacy: a utility-preserving formulation of differential privacy guarantees. *IEEE Transactions on Information Forensics and Security*, 12:1418-1429, 2017.